CLAIMS:

1. A method for reporting topology changes in a subnet of a switched fabric including at least a client, a subnet manager (SM) and switches interconnected via links, said method comprising:

creating and reporting a list of topology changes that are interesting to the client for topology change notifications;

when a topology change occurs in the subnet, determining if the topology change is in the list of topology changes created by the interested client; and

if the topology change is in the list of topology changes created by the interested client, reporting a topology change event to the interested client.

2. The method as claimed in claim 1, wherein said list of topology changes is created by the client to serve as client-defined filters that specify the types of topology changes the client is interested in receiving notifications.

3. The method as claimed in claim 2, wherein said list of topology changes includes, but is not limited to, when a new data path is created between a pair of end nodes in the subnet, when an existing data path is destroyed between a pair of end nodes in the subnet, when a new device is inserted in the subnet, and when an existing device is removed from the subnet.

4. The method as claimed in claim 1, wherein said client corresponds to an end node of the subnet having at least one channel adapter (CA) installed to support one or more ports for data communication via said links of the subnet.

5. The method as claimed in claim 2, wherein said determining the topology change in the list of topology changes and said reporting the topology change events to the interested client are executed by said subnet manager.

6. The method as claimed in claim 5, wherein said subnet manager (SM) is installed in another end node of the subnet, and is implemented either in hardware or software to provide management services for all switches and end nodes in the subnet.

7. The method as claimed in claim 5, wherein said subnet manager (SM) is installed in another end node of the subnet, and is implemented in software written using a high-level computer programming language for performing network management functions in compliance with the InfiniBand™ Architecture specification.

8. The method as claimed in claim 5, wherein said subnet manager (SM) is installed in another end node of the subnet for discovering the subnet topology, assigning unique

addresses to all ports that are connected to the subnet, and establishing possible data paths among all ports by programing switch forwarding tables for download to the switches in the subnet for routing data packets to destinations via possible data paths established between switch pairs.

9. The method as claimed in claim 1, wherein said client sends a VendorSet (SetNotificationFilter) message to the subnet manager (SM) after the list of topology changes is created to indicate the topology changes that require client notifications, and said subnet manager (SM) sends a VendorGetResp (SetNotificationFilter) message back to the interested client to confirm receipt of the list of topology changes that the client is interested.

10. The method as claimed in claim 1, wherein said subnet manager (SM) sends a VendorSend (TopologyChangeNotification) message to the interested client after the topology change is determined in the list of topology changes to notify the topology change that occurred, and said client sends a VendorSendResp (TopologyChangeNotification) message back to the subnet manager (SM) to acknowledge the topology change notification.

11. A data network, comprising:

a host system having at least one channel adapter (CA) installed therein supporting one or more ports;

at least one target system having at least one channel adapter (CA) installed therein

supporting one or more ports;

a switched fabric comprising a plurality of different switches which interconnect said host system via CA ports to said remote system via CA port along different physical links for data communications; and

a fabric manager provided in said host system for making topology discovery, assigning local identifiers (LIDs) to all ports that are connected in the switched fabric, and programming forwarding tables for switches in the switched fabric, wherein said fabric manager includes a topology change notification mechanism configured to provide topology change notifications by:

enabling a client at one of the host system and the target system to create and communicate a list of topology changes that are interesting to the client for topology change notifications;

determining if a topology change occurred in the switched fabric is in the list of topology changes created by the interested client; and

if the topology change is in the list of topology changes created by the interested client, reporting a topology change event to the interested client.

12. The data network as claimed in claim 11, wherein said list of topology changes is created by the client to serve as client-defined filters that specify the types of topology changes the client is interested in receiving topology change notifications.

13. The data network as claimed in claim 12, wherein said list of topology changes includes, but is not limited to, when a new data path is created between a pair of end nodes in the switched fabric, when an existing data path is destroyed between a pair of end nodes in the switched fabric, when a new device is inserted in the switched fabric, and when an existing device is removed from the switched fabric.

14. The data network as claimed in claim 11, wherein said fabric manager is installed in another one of the host system and the target system, and is implemented either in hardware or software to provide management services for all switches and end nodes in the switched fabric.

15. The data network as claimed in claim 11, wherein said fabric manager is installed in another one of the host system and the target system, and is implemented in software written using a high-level computer programming language for performing network management functions in compliance with the InfiniBand™ Architecture specification.

16. The data network as claimed in claim 15, wherein said fabric manager is further configured to discover the fabric topology, assign unique addresses to all ports that are connected to the switched fabric, and establish possible data paths among all ports by programing switch forwarding tables for download to the switches in the switched fabric for routing data packets to destinations via possible data paths established between switch pairs.

17. The data network as claimed in claim 11, wherein said client sends a VendorSet (SetNotificationFilter) message to the fabric manager after the list of topology changes is created to indicate the topology changes that require client notifications, and said fabric manager sends a VendorGetResp (SetNotificationFilter) message back to the interested client to confirm receipt of the list of topology changes that the client is interested.

18. The data network as claimed in claim 11, wherein said fabric manager sends a VendorSend (TopologyChangeNotification) message to the interested client after the topology change is determined in the list of topology changes to notify the topology change that occurred, and said client sends a VendorSendResp (TopologyChangeNotification) message back to the fabric manager to acknowledge the topology change notification.

19. A computer readable medium comprising instructions that, when executed by a host system in a switched fabric including end nodes and switches interconnected via links, cause the host system to:

enabling a client at an end node to create and communicate a list of topology changes that are interesting to the client for topology change notifications;

determining if a topology change occurred in the switched fabric is in the list of topology changes created by the interested client; and

if the topology change is in the list of topology changes created by the interested client, reporting a topology change event to the interested client.

20. The computer readable medium as claimed in claim 19, wherein said list of topology changes is created by the client to serve as client-defined filters that specify the types of topology changes the client is interested in receiving topology change notifications.

21. The computer readable medium as claimed in claim 20, wherein said list of topology changes includes, but is not limited to, when a new data path is created between a pair of end nodes in the switched fabric, when an existing data path is destroyed between a pair of end nodes in the switched fabric, when a new device is inserted in the switched fabric, and when an existing device is removed from the switched fabric.

22. The computer readable medium as claimed in claim 19, further causing the system to enable the client to send a VendorGetResp (SetNotificationFilter) message to the interested client upon receipt of a VendorSet (SetNotificationFilter) message from the interested client to confirm receipt of the list of topology changes that the client is interested.

23. The computer readable medium as claimed in claim 19, further causing the system to send a VendorSend (TopologyChangeNotification) message to the interested client after the

topology change is determined in the list of topology changes to notify the topology change that occurred, and to acknowledge the topology change notification upon receipt of a VendorSendResp (TopologyChangeNotification) message from the interested client.